arXiv:1604.07177v1 [stat.CO] 25 Apr 2016

On the Use of Penalty MCMC for Differential Privacy

Sinan Yıldırım

Faculty of Engineering and Natural Sciences
Sabancı University, İstanbul, Turkey

April 26, 2016

Abstract

We view the penalty algorithm of Ceperley and Dewing (1999), a Markov chain Monte Carlo (MCMC) algorithm for Bayesian inference, in the context of data privacy. Specifically, we study differential privacy of the penalty algorithm and advocate its use for data privacy. We show that in the simple model of independent observations the algorithm has desirable convergence and privacy properties that scale with data size. Two special cases are also investigated and privacy preserving schemes are proposed for those cases: (i) Data are distributed among several data owners who are interested in the inference of a common parameter while preserving their data privacy. (ii) The data likelihood belongs to an exponential family.

1 Introduction

One concern about sharing datasets for use of third parties is violation of privacy of individuals who contribute to those datasets with their private information. Significant amount of research is devoted to address that concern by either analysing privacy preserving properties of already existing methods or developing novel methods that are useful for statistical inference and still preserve data privacy to a quantifiable extent at the same time.

One popular formulation of data privacy is *differential privacy* (Dwork et al., 2006). Differential privacy associates a random algorithm with a mapping from the space of data to a space of probability distributions for the output of the algorithm. It then quantifies the privacy of the algorithm roughly by some 'difference' between two probability distributions corresponding to two 'neighbour' data sets that differ by only one entry. This 'difference' is expressed by two nonnegative parameters which are typically shown by ϵ and δ, and the algorithm is said to be (ϵ, δ)-differentially private. Specifically, let $\mathcal{A}$ be a random algorithm that takes input from a set $\mathcal{X} \subseteq \mathbb{R}^n$ for some $n \geq 1$ and produces random outputs in $\mathcal{S}$. Let $h(X, Y)$ be the edit distance or *Hamming distance* between data sets $X, Y \in \mathcal{X}$ i.e. number of entry-wise differences between X and Y.

Definition 1. Differential privacy: *We call a randomised algorithm $\mathcal{A}$ with domain $\mathcal{X}$ and range $\mathcal{S}$ (ϵ, δ)-differentially private if for any measurable set $S \subset \mathcal{S}$ and for all $X, Y \in \mathcal{X}$ such that $h(X, Y) \leq 1$, we*

have

$$P(\mathcal{A}(X) \in S) \leq \exp(\epsilon) P(\mathcal{A}(Y) \in S) + \delta.$$

where $\mathcal{A}(X)$ and $\mathcal{A}(Y)$ denotes random outputs of the algorithm using inputs X and Y, respectively. If $\delta = 0$, then $\mathcal{A}$ is called ϵ-differential private.

Obviously, the smaller ϵ and δ are the better privacy is preserved, since smaller ϵ and δ correspond to a closer match between the probability distributions of $\mathcal{A}(X)$ and $\mathcal{A}(Y)$.

1.1 Relevance of Monte Carlo methods

In Bayesian inference, the answer to an unknown parameter of interest is in the form of a probability distribution, namely the posterior distribution, denoted by π throughout the paper. Monte Carlo methods, which constitute a great proportion of Bayesian inference methods, are used to generate random samples either exactly or approximately from π when π is intractable, i.e. it does not have tractable expectations of certain functions of the parameter of interest. Therefore, by its nature Monte Carlo methods produce random outputs and that makes them natural candidates for privacy preserving methods.

Arguably the most widely used Monte Carlo method for Bayesian inference is Markov chain Monte Carlo (MCMC) (see Robert and Casella (2004) for example), where one generates a sequence of samples from a Markov chain that has π as its stationary distribution. Study of MCMC methods in relation to differential privacy has recently attracted researchers. For example, Wang et al. (2015) presented a modified version of the stochastic gradient Langevin dynamics algorithm of Welling and Teh (2011) for differential privacy. Even more recently, Foulds et al. (2016) discussed differential privacy of several MCMC methods.

1.2 Contribution

In this paper, we contribute to the literature on privacy preserving methods by studying differential privacy of a specific MCMC algorithm, namely the *penalty algorithm* (Ceperley and Dewing, 1999). The penalty algorithm is actually an approximation of the Metropolis-Hastings (MH) algorithm (Hastings, 1970), one of the most popular MCMC methods. The approximation is due to the use of the log-acceptance ratio of the MH algorithm in normally distributed noise with known variance.

A true implementation of the penalty algorithm requires being able to compute the log-acceptance ratio in the first place. That is why the penalty algorithm has so far been used with approximately normally distributed estimates of the log-acceptance ratio *only* when this ratio is intractable. Indeed, without privacy concerns, it is always favourable to use the log-acceptance ratio rather than its noisy version.

However, in a data privacy setting, the penalty algorithm becomes more relevant and potentially more useful than the MH algorithm, an observation that has surprisingly escaped attention so far. We advocate the use of the penalty algorithm for a few reasons:

- The penalty algorithm is an *exact* algorithm in that its stationary distribution remains π.
- The penalty algorithm deliberately uses noisy functions of data, hence it has desirable privacy preserving properties.

- Under certain conditions, the differential privacy as well as convergence properties of the penalty algorithm scale well with data size.

The first reason is already known (Ceperley and Dewing, 1999; Nicholls et al., 2012); we will justify the last two in Sections 2 and 3. Specifically, we first show that the penalty algorithm, although inferior to the MH algorithm that it mimics, can preserve the uniform or geometric ergodicity properties (with slower rates) of the MH algorithm under some conditions. Then we proceed with the analysis of the differential privacy of the algorithm and show that the algorithm does have desired privacy properties that scale favourably with data size. We also investigate the use of the algorithm in a *data sharing* scenario when data are distributed among different owners who want to infer the same θ conditional on all the available data .

Our analysis is similar to that in Wang et al. (2015) in that our main theorem (Theorem 1) employs the results on the so called *advanced composition* (Dwork et al., 2010) and *Gaussian mechanism* (Dwork and Roth, 2013). One difference is that the algorithm discussed in Wang et al. (2015) is only asymptotically exact but computationally scalable in data size due to subsampling; whereas our algorithm preserves the target density at all iterations but computationally not scalable in data size (unless the posterior distribution has a convenient form such as the exponential family, which we will discuss further in Section 3.3).

In a recent work by Foulds et al. (2016), differential privacy of certain MCMC methods for exponential family distributions are studied when they are used with the sufficient statistics of the data modified by Laplacian noise. The authors show that the proposed methods are asymptotically unbiased in terms of their target distribution. However, for any finite data size, the algorithms have bias. The penalty method, on the other hand, remains unbiased and it is applicable to more general distributions than exponential families.

1.3 Outline

The rest of the paper is organised as follows: In Section 2, we introduce the penalty algorithm for a general target distribution and we provide some statistical properties of the algorithm mainly in terms of convergence. The purpose of Section 2 is mainly to assist the analysis in Section 3 where we work out the differential privacy of the algorithm, specifically we will explain what happens to ϵ and δ of the algorithm as the data size increases. We conclude in Section 4 with final discussions and a mention to future work.

2 The penalty algorithm

Suppose we are given a target distribution π for variable $\theta \in \Theta \subseteq \mathbb{R}^{d_\theta}$ for some $d_\theta \geq 1$ with probability density shown by $\pi(\theta)$ and a proposal kernel q with density shown as $q(\cdot|\theta)$ for $\theta \in \Theta$. Under some mild conditions on q with respect to π, the Metropolis-Hastings (MH) algorithm (Hastings, 1970) generates a reversible Markov chain $\{\theta_t\}_{t\geq 0}$ with invariant density π as follows: Start with θ_0; for $t \geq 0$, given $\theta_t = \theta$, propose $\theta' \sim q(\cdot|\theta)$ and set $\theta_{t+1} = \theta'$ (or 'accept' θ') with probability

$$\alpha(\theta, \theta') = \min\left\{1, r(\theta, \theta')\right\},$$

else set $\theta_{t+1} = \theta'$ (or 'reject'), where

$$r(\theta, \theta') = \frac{\pi(\theta')q(\theta|\theta')}{\pi(\theta)q(\theta'|\theta)},$$

is called the acceptance ratio given the current and proposed values θ and θ'.

The penalty algorithm of Ceperley and Dewing (1999) replaces the logarithm of the acceptance ratio $\lambda(\theta, \theta') := \log r(\theta, \theta')$ with an unbiased normal estimate $\hat{\lambda}_\sigma(\theta, \theta')$ with variance $\sigma^2(\theta, \theta') = \sigma^2(\theta', \theta)$ that takes its values from a symmetric function $\sigma^2 : \Theta^2 \to [0, \infty)$, i.e.

$$\hat{\lambda}_\sigma(\theta, \theta') \sim \mathcal{N}(\lambda(\theta, \theta'), \sigma^2(\theta, \theta'))$$

and accepts θ' with the modified acceptance probability

$$\hat{\alpha}_\sigma(\theta, \theta') = \min\{1, \hat{r}_\sigma(\theta, \theta')\}$$

where the modified acceptance ratio

$$\hat{r}_\sigma(\theta, \theta') = \exp\left(\hat{\lambda}_\sigma(\theta, \theta') - \sigma^2(\theta, \theta')/2\right)$$

has the *penalty term* $\sigma^2(\theta, \theta')/2$ as a compromise for using an estimate $\hat{\lambda}_\sigma(\theta, \theta')$ instead of $\lambda(\theta, \theta')$.

2.1 Statistical properties

The penalty algorithm has the remarkable property that the Markov chain it generates still has π as its invariant density, thus the algorithm still targets π; see Ceperley and Dewing (1999) for the original proof and Nicholls et al. (2012) for an alternative proof that shows that the penalty algorithm is reversible with respect to π.

Nicholls et al. (2012) also showed (in a more general setting) that using $\hat{\alpha}_\sigma(\theta, \theta')$ is always inferior to using $\alpha(\theta, \theta')$ in the Peskun sense, i.e. for any θ, θ' the expected acceptance probability

$$\alpha_\sigma(\theta, \theta') := \mathbb{E}\left[\hat{\alpha}_\sigma(\theta, \theta')\right] \leq \alpha(\theta, \theta'),$$

which can be verified via Jensen's inequality. In addition, we show here that in fact the algorithm's performance in the Peskun sense (i.e. $\alpha_\sigma(\theta, \theta')$) decreases monotonically in $\sigma^2(\theta, \theta')$. (A simple proof by working out the derivative of $\alpha_\sigma(\theta, \theta')$ is given in the Appendix.)

Proposition 1. *For any $\theta, \theta' \in \Theta$, $\alpha_\sigma(\theta, \theta')$ is a decreasing function of $\sigma^2(\theta, \theta') \in [0, \infty)$.*

Denote the transition kernel of the MH algorithm and the corresponding penalty algorithm with variance $\sigma^2(\theta, \theta')$ by T and T_σ, respectively. Despite its (perhaps not surprising) inferiorities that are mentioned above, T_σ inherits the favourable ergodicity properties of the T under some conditions. The following proposition on uniform ergodicity is proven in a more general setting in Nicholls et al. (2012); we present it here specifically for the penalty algorithm for completeness.

Proposition 2. *Suppose that $\sup_{\theta,\theta'} \sigma^2(\theta, \theta') = B < \infty$. If the MH algorithm that targets π with proposal q is uniformly ergodic, then the penalty algorithm that uses $\sigma^2(\theta, \theta')$ for the variance of its noise and the same q for the proposal is also uniformly ergodic.*

Proof. We need to show that the expected acceptance probability of the penalty algorithm satisfies

$$\alpha_\sigma(\theta, \theta') \geq \kappa \alpha(\theta, \theta') \tag{1}$$

for some $\kappa > 0$ for all θ, θ'. Then we can follow the same steps starting from Equation (6.3) of Nicholls et al. (2012, Appendix A) to conclude that for any $\theta \in \Theta$ and measurable set $E \subseteq \Theta$ we have $T_\sigma(\theta, E) \geq \kappa T(\theta, E)$ and that since by hypothesis T satisfies a minorisation condition so does T_σ, which leads to uniform ergodicity. To show (1), let $\kappa = 1 - \Phi(B/2)$, where Φ is the cumulative distribution function of $\mathcal{N}(0, 1)$, and observe that, for any $\theta, \theta' \in \Theta$, and $V \sim \mathcal{N}(0, 1)$, we have

$$\begin{aligned} P\left(\hat{\lambda}_\sigma(\theta, \theta') \geq \lambda(\theta, \theta')\right) &= P\left(\sigma(\theta, \theta')V - \sigma^2(\theta, \theta')/2\right) \\ &= P\left(V \geq \sigma(\theta, \theta')/2\right) \\ &\geq P\left(V \geq B/2\right) \\ &= 1 - \Phi(B/2) = \kappa. \end{aligned}$$

Therefore, by using a decomposition for the expected value, we have

$$\begin{aligned} \alpha_\sigma(\theta, \theta') &= \mathbb{E}\left[\hat{\alpha}(\theta, \theta')\right] \\ &\geq \mathbb{E}\left[\hat{\alpha}_\sigma(\theta, \theta') \,\middle|\, \hat{\lambda}_\sigma(\theta, \theta') \geq \lambda(\theta, \theta')\right] P\left(\hat{\lambda}_\sigma(\theta, \theta') \geq \lambda(\theta, \theta')\right) \\ &\geq \alpha(\theta, \theta')\kappa \end{aligned}$$

where the second inequality holds since $\hat{\alpha}_\sigma(\theta, \theta') \geq \alpha(\theta, \theta')$ in the event $\hat{\lambda}_\sigma(\theta, \theta') \geq \lambda(\theta, \theta')$. □

Assumption of uniform ergodicity for the MH algorithm can be restrictive. Thus, whether the penalty algorithm inherits the geometric ergodicity of the MH algorithm (with a slower rate) is worth investigating. The following lemma, which uses an intermediate result of Proposition 2, will be useful to establish geometric ergodicity of the penalty algorithm.

Lemma 1. *Suppose the MH algorithm targeting π on Θ has the rejection probability function $\rho(\theta)$. Also, suppose the penalty algorithm for π that uses the same proposal kernel and variance $\sigma^2 : \Theta^2 \to [0, \infty)$ has the rejection probability function $\rho_\sigma(x)$. If $\operatorname{ess\,sup} \rho < 1$, the essential supremum being taken with respect to π, and $\sup_{\theta,\theta'} \sigma^2(\theta, \theta') = B < \infty$, then $\operatorname{ess\,sup} \rho_\sigma < 1$.*

Proof. We have shown in Proposition 2 that if $\sup_{\theta,\theta'} \sigma^2(\theta, \theta') = B < \infty$ then $\alpha_\sigma(\theta, \theta') \geq \kappa \alpha(\theta, \theta')$ with $\kappa = 1 - \Phi(B/2)$.

$$\begin{aligned} \rho_\sigma(\theta) &= 1 - \int \alpha_\sigma(\theta, \theta') q(d\theta'|\theta) \\ &\leq 1 - \kappa \int \alpha(\theta, \theta') q(d\theta'|\theta) \\ &= 1 - \kappa(1 - \rho(\theta)) \\ &= (1 - \kappa) + \kappa\rho(\theta) \end{aligned}$$

Therefore, $\operatorname{ess\,sup} \rho_\sigma \leq (1 - \kappa) + \kappa \operatorname{ess\,sup} \rho < 1$. □

Define the sub-stochastic operator $T_{\sigma,a}$ the acceptance part of the kernel for the penalty algorithm, i.e. for any $\varphi \in L^2(\pi)$

$$T_{\sigma,a}\varphi(\theta) = \int q(d\theta'|\theta)\alpha_\sigma(\theta,\theta')\varphi(\theta')$$

We give a result on the geometric ergodicity of the penalty method that holds under the assumption that $T_{\sigma,a}$ is compact. (See Atchadé and Perron (2007) for a discussion on the compactness assumption.)

Proposition 3. *Assume $T_{\sigma,a}$ is compact. If the MH algorithm targeting π is geometrically ergodic, then the penalty algorithm using a bounded $\sigma^2(\theta,\theta')$ is also geometrically ergodic.*

Proof. The geometric ergodicity of the MH algorithm implies $\operatorname{ess\,sup}\rho < 1$ by Roberts and Tweedie (1996, Proposition 5.1.). The boundedness of $\sigma^2(\theta,\theta')$ gives $\operatorname{ess\,sup}\rho_\sigma < 1$ by Lemma 1. Moreover, we observe that T_σ is self-adjoint since the penalty algorithm is reversible. These facts can be used to establish the geometric ergodicity by verifying the steps of Theorem 2.1. of Atchadé and Perron (2007) for the kernel of the penalty algorithm. □

3 Differential privacy of the penalty algorithm

We have mentioned above that the MH algorithm is superior to the penalty algorithm that mimics it with an extra noise term. Then, what is the point of using the penalty algorithm in the first place? There are at least two reasons, the second of which is relevant to data privacy.

1. $\lambda(\theta,\theta')$ may be intractable or too expensive to compute and instead an *approximately* normal estimate of it may be used (in which case the resulting algorithm becomes an *approximate* penalty algorithm in that the target density is no more π). This is in fact how the algorithm has been used in the literature so far.

2. Adding noise to $\lambda(\theta,\theta')$ may help with preserving some sort of *data privacy* in a Bayesian framework where π, hence $\lambda(\theta,\theta')$, depends on the data to be conditioned on.

The second argument is the main motivation of this paper and needs elaboration with a model that we will study for the rest of the paper. Suppose that we are interested in Bayesian inference of θ where π is the posterior distribution given the data $y_{1:n}$ of size n with prior distribution having a probability density function η,

$$\pi(\theta) := p(\theta|y_{1:n}) \propto \eta(\theta)p(y_{1:n}|\theta),$$

and the data samples conditional on θ are independent, which leads to the following simple factorisation of the data likelihood

$$p(y_{1:n}|\theta) = \prod_{t=1}^{n} p(y_t|\theta).$$

Let us rewrite the logarithm of the acceptance ratio as $\lambda_n(y_{1:n},\theta,\theta') = d_n(y_{1:n},\theta,\theta') + \log\frac{\eta(\theta')q(\theta|\theta')}{\eta(\theta)q(\theta'|\theta)}$ where

$$d_n(y_{1:n},\theta,\theta') := \sum_{t=1}^{n}\left[\log p(y_t|\theta') - \log p(y_t|\theta)\right]$$

where we include n as well as $y_{1:n}$ in the notation to indicate dependency on data. Assuming calculation of $\log \frac{\eta(\theta')q(\theta|\theta')}{\eta(\theta)q(\theta'|\theta)}$ is straightforward, in order to implement the penalty algorithm, one only needs a normally distributed noisy version of $d_n(y_{1:n}, \theta, \theta')$ provided that the variance of the noise is known. Therefore, the data owner can feed the analyser with the noisy log-acceptance ratio and ensure a certain degree of data privacy as well as allowing analysis of the data. Another relevant scenario is where the data are *shared* among a certain number of users whose common interest is inferring θ and common concern is privacy of their data against each other. As we will discuss later, the penalty algorithm suggests that those users can execute a common algorithm by submitting their contributions to the log acceptance ratio in gaussian noise with known variance.

The important questions to be asked in the scenarios considered above are:

- How should we choose the noise variance $\sigma_n^2(\theta, \theta')$ to ensure a certain degree of data privacy?
- How does the differential privacy of such schemes scale with data size n?

These questions will be addressed in the rest of this section, mainly by Theorem 1. In the following, we will investigate the degree to which differential privacy is ensured by enabling implementation of the penalty algorithm with reasonable convergence properties. We will first give our main result on the differential privacy of the penalty algorithm. We then provide an interpretation of the result and comment about the assumptions. Then we will provide a proof of the main result which relies on the concepts of *advanced composition* (Dwork et al., 2010) and *Gaussian mechanism* (Dwork and Roth, 2013).

3.1 Reviewing the penalty algorithm in the privacy context

In order to study the differential privacy of the penalty algorithm, it is useful to view its iterations as a sequence of *database access mechanisms* where these mechanisms are (randomly) called upon outside by an *adversary*. In particular, we need to introduce the concept of the *Gaussian mechanism* (Dwork and Roth, 2013). Define the Gaussian mechanism with $\sigma^2 > 0$ for a function $f : \mathcal{X} \to \mathbb{R}^d$ with some $d \geq 1$ to be an algorithm $\mathcal{A}_\sigma : \mathcal{X} \to \mathbb{R}^d$ which outputs a noisy version of $f(X)$ for input X:

$$\mathcal{A}_\sigma(X) = \hat{f}(X) \sim \mathcal{N}(f(X), \sigma^2 I_d).$$

For the penalty algorithm, for $\theta, \theta' \in \Theta$, $n \geq 1$ and a non-negative real number $\sigma > 0$, we define the database access mechanism $\mathcal{M}_{\theta,\theta',n,\sigma}$ to be the Gaussian mechanism with variance σ^2 for the function $d_n(\cdot, \theta, \theta') : \mathcal{Y}^n \to \mathbb{R}$. That is, given the data $y_{1:n}$, $\mathcal{M}_{\theta,\theta',n,\sigma}$ returns

$$\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta') \sim \mathcal{N}(d_n(y_{1:n}, \theta, \theta'), \sigma^2). \tag{2}$$

Equation (2) is nothing but the noisy version of the likelihood ratio in the penalty algorithm.

Given the data $y_{1:n}$, initial value $\theta_0 \in \Theta$ and variance $\sigma_n^2(\theta, \theta') = \sigma^2 > 0$ for all $\theta, \theta' \in \Theta$, let the sequence of generated samples and proposed samples of the penalty algorithm be $\{\theta_{n,t}, t \geq 1\}$ and $\{\theta'_{n,t}, t \geq 1\}$ respectively. Given the current sample $\theta_{n,t-1} = \theta$ and the proposed sample $\theta'_{n,t} = \theta'$, the next iteration of the penalty algorithm can be viewed as in Algorithm 1:

Algorithm 1: Iteration t of the penalty method with a database access view

Input: $\theta_{n,t-1} = \theta$, $\theta'_{n,t} = \theta'$

Output: $\theta_{n,t}$, $\theta'_{n,t+1}$, $\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta')$

1 The database access mechanism $\mathcal{M}_{\theta,\theta',n,\sigma}$ returns $\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta')$ as in (2).

2 The adversary decides on $\theta_{n,t}$ according to the acceptance probability

$$\min\left\{1, \frac{\eta(\theta')q_n(\theta|\theta')}{\eta(\theta)q_n(\theta'|\theta)} \exp\left(\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta') - \sigma^2/2\right)\right\}.$$

The adversary then samples $\theta'_{n,t+1}$ from $q_n(\cdot|\theta_{n,t})$.

The *view of the adversary* using the penalty algorithm targeting π_n for k iterations is

$$\{\theta_{n,t}, \theta'_{n,t}, \hat{d}_{n,\sigma}(y_{1:n}, \theta_{n,t-1}, \theta'_{n,t}); 1 \leq t \leq k\}.$$

That is, we consider the scenario in which not only the generated samples but also the estimates of the log acceptance ratios are observed during the run of the penalty algorithm. For simplicity of the analysis and emphasis on the contribution of the accept-reject procedure of the penalty algorithm, we will from here on assume that the proposal distribution q_n does not depend on $y_{1:n}$ (more on this assumption later). Under that assumption, observe that only Step 1 of the above description requires access to data. Therefore, the differential privacy of the penalty algorithm lies in the differential privacy of a sequence the Gaussian mechanisms called during the iterations.

3.2 The main result

Denote the joint probability distribution of sampled and proposed variables by P_{n,σ,θ_0} ($y_{1:n}$ is omitted from the notation for simplicity) so that for every $k \geq 1$

$$P_{n,\sigma,\theta_0}(d\theta_{n,1:k} \times d\theta'_{n,1:k}) = \prod_{t=1}^{k} q_n(d\theta'_{n,t}|\theta_{n,t-1}) \left[\alpha_\sigma(\theta_{n,t-1}, \theta'_{n,t})\delta_{\theta'_{n,t}}(d\theta_{n,t}) + (1 - \alpha_\sigma(\theta_{n,t-1}, \theta'_{n,t}))\delta_{\theta_{n,t-1}}(d\theta_{n,t})\right]$$

Our main theorem requires the following assumption on the structure of the log-likelihood together with the choice of the proposal density to hold P_{n,σ,θ_0}-a.s.

Assumption A1. *For all $\theta \in \Theta$, there exists an $\alpha > 0$ such that $\nabla_2 d_n(\cdot, \theta, \theta') \leq cn^{-\alpha}$ P_{n,σ,θ_0}-a.s., where the ℓ_2 sensitivity $\nabla_2 f$ of the function $f : \mathcal{X} \to \mathbb{R}^d$ for some $d \geq 1$ is defined as*

$$\nabla_2 f = \sup_{x,y \in \mathcal{X}: h(x,y) \leq 1} ||f(x) - f(y)||_2.$$

Theorem 1 states the number of iterations allowed to have a (ϵ, δ) differential privacy with its ϵ and δ parameters having desirable orders of magnitude, which are $\mathcal{O}(1)$ and $o(1/n)$, respectively (see the discussion in Dwork and Roth (2013, Section 2.3)).

Theorem 1. *Assume A1 with some $c > 0$ and $\alpha > 0$. For every $\beta > 0$, $k_0 > 0$ and $\theta_0 \in \Theta$, there exists some $\sigma^2 > 0$ such that it holds P_{n,σ,θ_0}-a.s. that the $k(n) = \lfloor k_0 n^{2\alpha}/\log(n) \rfloor$ iterations of the penalty algorithm targeting π_n and using $\sigma_n^2(\theta, \theta') = \sigma^2$ is $(\epsilon_{n,k(n)}, \delta_{n,k(n)})$-differentially private where*

$$\epsilon_{n,k(n)} \leq 2\sqrt{k_0(2\alpha+\beta)\beta}\frac{c}{\sigma} + 4k_0\frac{c^2}{\sigma^2}(2\alpha+\beta) \tag{3}$$

$$\delta_{n,k(n)} \leq 1.25k_0 n^{-\beta}/\log(n) + n^{-\beta}. \tag{4}$$

In particular, for $\beta > 1$, we have $\epsilon_{n,k(n)} = \mathcal{O}(1)$ and $\delta_{n,k(n)} = o(1/n)$.

Remarks: Theorem 1 is useful due to its explicit expressions for $k(n)$, $\epsilon_{n,k(n)}$, and $\delta_{n,k(n)}$. (It is possible to read (3) and (4) as equalities, since if an algorithm is (ϵ, δ) differentially private, it is also (ϵ', δ') differentially private for $\epsilon' > \epsilon$ and $\delta' > \delta$.) Specifically, it provides the order of how many iterations the penalty algorithm should be run and what σ^2 should be taken to achieve a given degree of differential privacy. In particular, we are interested in $\epsilon_{n,k(n)} = \mathcal{O}(1)$ and $\delta_{n,k(n)} = o(1/n)$ since those are the desired orders for the algorithm to be meaningfully private, see the discussion in Dwork and Roth (2013, Section 2.3).

What the theorem suggests is that by tuning the free parameters k_0 and β we can obtain a 'scalable' differential privacy for a non-decreasing (in fact, increasing for large n) number of iterations of the penalty algorithm whose variance does not increase with n. Note that while α and c depend on the model at hand, we have the freedom choose k_0, β. A realistic value for α is 0.5, as we will discuss in Section 3.2.1 also.

The depicted scenario in which we observe the noisy log-likelihood ratio $\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta')$ will be more relevant when we discuss the use of the penalty algorithm in a data sharing scenario. In the absence of such a case, one can consider instead a scenario where the view of the adversary is only the posterior samples and the proposed samples $\theta_{n,t}, \theta'_{n,t}$, $t \geq 1$, or even only the posterior samples $\theta_{n,t}$, $t \geq 1$. In such a scenario, we can still use $(\epsilon_{n,k(n)}, \delta_{n,k(n)})$ to describe the differential privacy of the such a scheme. This is because once $\hat{d}_{n,\sigma}(y_{1:n}, \theta, \theta')$ is calculated, the accept-reject step and the proposal step can be seen as random *post-processing steps* that do not require any access to data. This can only result in an algorithm whose differential privacy after $k(n)$ iterations is less than or equal to that $(\epsilon_{n,k(n)}, \delta_{n,k(n)})$ given in (3) and (4).

Proof of the main result: The following theorem by Dwork and Roth (2013) on the Gaussian mechanism plays a significant role in establishing the proof of Theorem 1.

Theorem 2. ***(Gaussian mechanism)*** *Let $\epsilon \in (0, 1)$ be arbitrary. The Gaussian mechanism for function f with $\sigma > \nabla_2 f\sqrt{2\log(1.25/\delta)}/\epsilon$ is (ϵ, δ)-differentially private.*

The function f relevant to the penalty algorithm is the $d_n(\cdot; \theta, \theta') : \mathcal{Y}^n \to \mathbb{R}$'s, i.e. the difference of log likelihood functions, when considered as functions of data samples. We need to cite one more theorem about differential privacy, regarding what happens when a sequence of database access mechanisms with a certain differential privacy is applied (Dwork et al., 2010).

Theorem 3. ***(Advanced composition)*** *For all $\epsilon, \delta, \delta' \geq 0$, the class of (ϵ, δ)-differentially private mechanisms satisfies $(\epsilon', k\delta + \delta')$-differential privacy under k-fold adaptive composition for:*

$$\epsilon' = \sqrt{2k\log(1/\delta')}\epsilon + k\epsilon(e^\epsilon - 1).$$

The structure of Algorithm 1 suggests that the penalty algorithm can be seen as a k-fold adaptive composition, where a database access mechanism with a certain differential privacy is called. This, combined with Theorems 2 and 3 facilitates the proof of Theorem 1.

Proof. (Theorem 1)
Let $\beta' = 2\alpha + \beta$ and pick a finite $\sigma > 0$ that satisfies

$$\sigma > \max_{n \geq 1} cn^{-\alpha}\sqrt{2\beta' \log n}. \tag{5}$$

which is possible since the right hand side is finite for each $\alpha > 0$. Next, let $\mathcal{D}_{c,\alpha,n,\sigma} = \{\mathcal{M}_{\theta,\theta',n,\sigma} : \nabla_2 d_n(\cdot, \theta, \theta') < cn^{-\alpha}\}$ be a family of the Gaussian mechanisms where $\mathcal{M}_{\theta,\theta',n,\sigma}$ was defined earlier. Then, Theorem 2 implies that each mechanism in $\mathcal{D}_{c,\alpha,n,\sigma}$ is (ϵ, δ)-DP where

$$\epsilon = \frac{c}{\sigma} n^{-\alpha}\sqrt{2\beta' \log n}, \quad \delta = 1.25n^{-\beta'}$$

since $\sigma > cn^{-\alpha}\sqrt{2\beta' \log n}$ by (5), hence satisfying the condition of Theorem 2.

Furthermore, by Theorem 3, we see that for $\delta'_n > 0$ and $k \geq 1$, $\mathcal{D}_{c,\alpha,n,\sigma}$ satisfies $(\epsilon_{n,k}, \delta_{n,k})$-DP under k-fold adaptive composition

$$\epsilon_{n,k} \leq \sqrt{4k\beta' \log n \log(1/\delta'_n)}\frac{c}{\sigma} n^{-\alpha} + 2k\frac{c^2}{\sigma^2} n^{-2\alpha} 2\beta' \log n, \quad \delta_{n,k} = 1.25kn^{-\beta'} + \delta'_n.$$

where we used $e^\epsilon - 1 \leq 2\epsilon$ since $\epsilon < 1$ (as noted by Wang et al. (2015) also). Now, substituting $\delta'_n = n^{-\beta}$ and $k = k(n) = \lfloor k_0 n^{2\alpha} / \log(n) \rfloor$, we have

$$\begin{aligned}
\epsilon_{n,k(n)} &\leq 2\sqrt{k_0 \beta' \beta}\frac{c}{\sigma} + 4k_0 \frac{c^2}{\sigma^2}\beta' \\
\delta_{n,k(n)} &\leq 1.25k_0 n^{-\beta} / \log(n) + n^{-\beta}.
\end{aligned}$$

We conclude the proof by substituting $\beta' = 2\alpha + \beta$ and pointing out that by Assumption A1 each iteration of the penalty algorithm targeting π_n accesses data via a mechanism from $\mathcal{D}_{c,\alpha,n,\sigma}$ P_{n,σ,θ_0}-a.s. □

3.2.1 On the assumptions about the proposal density

Although Assumption A1 is not easy to verify, the following two assumptions are more testable and they together imply A1.

Assumption A2. *The derivative with respect to θ of the log-likelihood $\ell(y, \theta)$ exists and for all $y \in \mathcal{Y}$ it is bounded by M in absolute value*

$$\sup_{y \in \mathcal{Y}, \theta \in \Theta} \left| \frac{\partial \ell(y, \theta)}{\partial \theta(i)} \right| \leq M, \quad i = 1, \ldots, d_\theta.$$

Assumption A3. *There exists an $\alpha > 0$ and $c > 0$ such that for all $n \geq 1$ we have $|\theta'_{n,t}(i) - \theta_{n,t-1}(i)| < cn^{-\alpha}$ P_{n,σ,θ_0}-a.s. for all $i = 1, \ldots, d_\theta$.*

Assumption A2 can hold for compact sets for θ and y and Assumption A3 can be forced, for example, with a truncated random walk. Assumptions A2 and A3 imply Assumption A1, since

$$|\ell(y,\theta') - \ell(y,\theta)| \leq \sum_{i=1}^{d_\theta} |\theta'(i) - \theta(i)| M$$

by the zeroth order Taylor polynomial approximation and therefore $\nabla_2 d_n(\cdot, \theta, \theta') \leq 2 d_\theta n^{-\alpha} M$. For Assumption A3 to be compatible with desirable convergence properties for the algorithm, we would need $\alpha \leq 0.5$, since for any $\alpha > 0.5$ the proposal moves would have too short of a range and this would cause the penalty algorithm to collapse as $n \to \infty$.

A restrictive assumption we made about q_n is its independence from the data $y_{1:n}$. This can be suboptimal when compared to schemes that exploiting the gradient of the posterior π_n that can only be extracted by using the data. We made the independence assumption merely for emphasising on the differential privacy brought by the Gaussian mechanism that is naturally a part of the penalty algorithm. When q_n depends on the data, sampling from q_n should be seen as a database access mechanism and therefore the analysis of Algorithm 1 should consider the differential privacy of the proposal step as well. Taking this into account, any proposal scheme whose differential privacy at one iteration is as good as the Gaussian mechanism is welcome, since such a scheme will not distort the order of $\epsilon_{n,k(n)}$ and $\delta_{n,k(n)}$ but will only double these quantities in the worst case.

3.2.2 Differential privacy of the penalty algorithm in a data sharing context

Consider a scenario where the available data $y_{1:n}$ is shared among N *data owners (sharers)* who are commonly interested in inferring the common parameter θ using all the data available. Let the portion of the data shared by sharer i be $y^{(i)}$ of size n_i, with $n_1 + \ldots + n_N = n$, so that $y_{1:n} = (y^{(1)}, \ldots, y^{(N)})$. Suppose these sharers want to implement a penalty algorithm together that targets π_n as well as guarantees a certain differential privacy for each sharer against the other sharers. In other words, for each sharer all the other sharers are adversaries. This is possible if each sharer contributes to the log-acceptance with added noise. Specifically, given $\theta, \theta' \in \Theta$ as the current and proposed values in an iteration of the penalty algorithm, the contribution of user i shall be

$$\hat{d}^{(i)}_{n,\sigma}(y^{(i)}, \theta, \theta') \sim \mathcal{N}(d_{n_i}(y^{(i)}, \theta, \theta'), \sigma^2). \tag{6}$$

The steps of this algorithm can be written as in Algorithm 2. The view of the adversaries (which are data

Algorithm 2: Iteration t of the penalty method within a data sharing context

1 For $i = 1, \ldots, N$, user i returns its contribution in noise, $\hat{d}^{(i)}_{n,\sigma}(y^{(i)}, \theta, \theta')$ as in (6).

2 $\theta_{n,t}$ is decided on according to the acceptance probability

$$\min\left\{1, \frac{\eta(\theta') q_n(\theta|\theta')}{\eta(\theta) q_n(\theta'|\theta)} \exp\left(\sum_{i=1}^{N} \hat{d}^{(i)}_{n,\sigma}(y^{(i)}, \theta, \theta') - N\sigma^2/2\right)\right\}.$$

Then the proposal $\theta'_{n,t+1}$ is sampled from $q_n(\cdot|\theta_{n,t})$ for the next iteration.

sharers at the same time) using the penalty algorithm targeting π_n for k iterations is

$$\{\theta_{n,t}, \theta'_{n,t}, \hat{d}^{(i)}_{n,\sigma}(y^{(i)}, \theta_{n,t-1}, \theta'_{n,t}), 1 \leq i \leq N, 1 \leq t \leq k\}.$$

Note that the ℓ_2 sensitivity of the function that each data sharer reveals in noise is the same as in Section 3.2. Therefore, with the same σ^2, the result for differential privacy in Theorem 1 applies for the differential privacy of each user. However, there is a cost to pay: The total variance variance in the log-acceptance ratio is now $N\sigma^2$ instead of σ^2 which slows down the mixing of the chain of the penalty algorithm and makes it worse in Peskun sense. However, if N does not increase with n, Algorithm 2 is viable since then the variance of the total noise in the algorithm will increase not with n.

3.3 Exponential families

A broad and useful family of distributions for the likelihood are exponential family distributions where the likelihood of the data given θ is written as

$$p(y|\theta) = h(y)g(\theta)\exp[\varphi(\theta)^T S(y)]$$

where $S(y)$ is a $d_\varphi \times 1$ vector of *sufficient statistics* of y, $\varphi(\theta)$ is the $d_\varphi \times 1$ vector of natural parameters of the model, and $h(y)$ is the normalising constant. For the i.i.d. model with n data samples, the logarithm of the ratio of log-likelihoods at θ and θ' becomes

$$d_n(y_{1:n}, \theta, \theta') = n[\log g(\theta') - \log g(\theta)] + [\varphi(\theta') - \varphi(\theta)]^T S_n(y_{1:n}) \tag{7}$$

where $S_n(y_{1:n}) = \sum_{i=1}^n S(y_i)$.

Obviously, the previously discussed settings and stated results also hold for exponential family distributions; however, exponential families enable an *alternative and simpler* way of sharing data in privacy: As one can see from (7), a noisy version of $d_n(y_{1:n}, \theta, \theta')$ can be provided by providing a *noisy version of the sufficient statistic* $S_n(y_{1:n})$. In order to implement the penalty algorithm for k iterations, k independent noisy versions of $S_n(y_{1:n})$ will be revealed by the data owner, i.e.

$$\hat{S}^{(t)}_{n,\xi}(y_{1:n}) \overset{\text{i.i.d.}}{\sim} \mathcal{N}(S_{n,\sigma}(y_{1:n}), \xi_n I_d), \quad t = 1, \ldots, k.$$

The resulting log-ratio of likelihood is then

$$\hat{d}_{n,\xi}(y_{1:n}, \theta, \theta') = n[\log g(\theta') - \log g(\theta)] + [\varphi(\theta') - \varphi(\theta)]^T \hat{S}_n(y_{1:n})$$

It can be checked that variance of $\hat{d}_{n,\xi}(y_{1:n}, \theta, \theta')$ can be written as

$$\sigma^2_n(\theta, \theta') = [\varphi(\theta') - \varphi(\theta)]^T [\varphi(\theta') - \varphi(\theta)]\xi_n. \tag{8}$$

The resulting penalty algorithm for exponential families is given in Algorithm 3. Note that unlike the method in Foulds et al. (2016) which adds noise to $S_n(y_{1:n})$ only once and then processes it, the penalty algorithm requires independent noisy sufficient statistics as many as the number of iterations it runs for. Another

Algorithm 3: Iteration t of the penalty method for exponential families

Input: $\theta_{n,t-1} = \theta$, $\theta'_{n,t} = \theta'$

Output: $\theta_{n,t}$, $\theta'_{n,t+1}$, $\hat{S}_{n,\xi}(y_{1:n})$

1 The database access mechanism returns and independent $\hat{S}_{n,\sigma}(y_{1:n}) \sim \mathcal{N}(S_{n,\xi}(y_{1:n}), \xi_n I_d)$ as in (2).

2 The adversary calculates $\sigma_n^2(\theta, \theta') = [\varphi(\theta') - \varphi(\theta)]^T [\varphi(\theta') - \varphi(\theta)]\xi_n$ and decides on $\theta_{n,t}$ according to the acceptance probability

$$\min\left\{1, \frac{\eta(\theta')q_n(\theta|\theta')}{\eta(\theta)q_n(\theta'|\theta)} \exp\left(\hat{d}_{n,\xi}(y_{1:n}, \theta, \theta') - \sigma_n^2(\theta, \theta')/2\right)\right\}.$$

The adversary then samples $\theta'_{n,t+1}$ from $q_n(\cdot|\theta_{n,t})$.

difference from the work of Foulds et al. (2016) is that their method is only asymptotically unbiased in terms of the target distribution whereas the penalty method is always unbiased.

There are two computation-wise important observations for Algorithm 3. Firstly, the database access mechanism is much simpler than the general case since it returns independent noisy versions of the same quantity. Secondly, unlike for general models, the computation load does not grow with n for exponential families.

Differential privacy of k iterations of Algorithm 3 can be computed using Theorems 2 and 3 once we know the l_2 sensitivity of S:

$$\nabla_2 S = \sup_{x,y \in \mathcal{Y}} ||S(x) - S(y)||_2$$

On one hand, due to convergence issues discussed in Section 2.1, it is important to keep $\sigma_n^2(\theta, \theta)$ non-increasing with n. On the other hand, $\sigma_n^2(\theta, \theta')$ should be large enough to satisfy differential privacy for sufficiently many iterations. Fortunately, Corollary 1 shows that we can do these two things at the same time under Assumption A4 which is essentially the counterpart of Assumption A1.

Assumption A4. *There exists an $\alpha > 0$ and $c > 0$ such that for all $n \geq 1$ we have $|\varphi(\theta'_{n,t}) - \varphi(\theta_{n,t-1})| < cn^{-\alpha}$ P_{n,σ,θ_0}-a.s.*

Again, counterparts of Assumptions A2 and A3 on φ and q_n can be made to verify Assumption A4, we do not give them here to avoid repetition.

Corollary 1. *Assume A4 with some $c > 0$ and $\alpha > 0$. For every $\beta > 0$, $k_0 > 0$ and $\theta_0 \in \Theta$, there exists some $\sigma^2 > 0$ such that it holds P_{n,σ,θ_0}-a.s. that the $k(n) = k_0 n^{2\alpha}/\log(n)$ iterations of Algorithm 3 targeting π_n and using $\xi_n = \sigma^2(\nabla_2 S)^2 n^{2\alpha}/c^2$ is $(\epsilon_{n,k(n)}, \delta_{n,k(n)})$-differentially private where $\epsilon_{n,k(n)}$ and $\delta_{n,k(n)}$ are given in (3) and (4), respectively. Moreover, the variance of the noise used in the algorithm is bounded: for all $n \geq 1$, $\theta, \theta' \in \Theta$ $\sigma_n^2(\theta, \theta') \leq (\nabla_2 S)^2 \sigma^2$.*

Proof. The proof is similar to that of Theorem 1. Let $\beta' = 2\alpha + \beta$ and pick a finite $\sigma > 0$ such that $\max_{n \geq 1} cn^{-\alpha}\sqrt{2\beta' \log n}$. The Gaussian mechanism that returns $\hat{S}_{n,\xi}$ with $\xi_n = \sigma^2(\nabla_2 S)^2 n^{2\alpha}/c^2$ is (ϵ, δ)-DP where $\epsilon = \frac{c}{\sigma} n^{-\alpha}\sqrt{2\beta' \log n}$ and $\delta = 1.25 n^{-\beta'}$ by Theorem 2. The rest of the proof for differential privacy is

the same as in the proof of Theorem 1. The bound for $\sigma_n^2(\theta, \theta')$ should be obvious from the choice made for ξ_n, equation (8) and Assumption A4. □

4 Conclusion

The main contribution of this paper arises from the simple observation that the penalty algorithm has a built-in noise in its calculations which is not desirable in any other context but can be exploited for data privacy. We have shown that, without letting the variance of the noise increase and hence convergence properties of the algorithm degrade with increasing data size, the differential privacy of the algorithm can scale well with data size. The penalty algorithm has the advantage over the other Bayesian methods for differential privacy in the sense that it is always unbiased in terms of its target distribution.

Our analysis uses the fact that at each iteration a database access mechanism with the same differential privacy is called upon. This is necessary for employing Theorem 3 on advanced composition. It would be worth studying what happens when mechanisms with different (but characterisable) differential privacies are called sequentially. This would help us relax our assumption on the bounded sensitivity of the difference of the log-likelihood.

We have already stated that our algorithm preserves the target density at all iterations but computationally not scalable in data size (unless the posterior distribution has a convenient form such as exponential families). A computationally scalable version of a privacy preserving penalty algorithm can be developed via subsampling, but this will introduce bias in the target density that may vanish asymptotically in data size under some conditions. This is in fact the topic of an unpublished work carried out by the author (in collaboration).

Acknowledgement

Thanks to Canan Pehlivan and Taylan Cemgil for their careful reading of this paper.

References

Atchadé, Y. F. and Perron, F. (2007). On the geometric ergodicity of Metropolis-Hastings algorithms. *Statistics*, 41(1):77–84.

Ceperley, D. M. and Dewing, M. (1999). The penalty method for random walks with uncertain energies. *Journal of Chemical Physics*, 110(20):9812–9820.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cyrptography*, pages 265–284. Springer.

Dwork, C. and Roth, A. (2013). The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407.

Dwork, C., Rothblum, G. N., and Vadhan, S. (2010). Boosting and differential privacy. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 51–60.

Foulds, J., Geumlek, J., and an Kamalika Chaudhuri, M. W. (2016). On the theory and practice of privacy-preserving Bayesian data analysis. Technical report, arxiv:1603.07294.

Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 52(1):97–109.

Nicholls, G., Fox, C., and Watt, A. (2012). Coupled MCMC with a randomised acceptance probability. Technical report, arXiv:1205.6857.

Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. New York: Springer, 2 edition.

Roberts, G. and Tweedie, R. (1996). Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83:95–110.

Wang, Y.-X., Fienberg, S., and Smola, A. (2015). Privacy for free: Posterior sampling and stochastic gradient Monte Carlo. In Blei, D. and Bach, F., editors, *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2493–2502. JMLR Workshop and Conference Proceedings.

Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In Getoor, L. and Scheffer, T., editors, *Proc. 28th International Conf. Machine Learning (ICML 2011)*, ICML '11, pages 681–688, New York, NY, USA. ACM.

A Appendix

Proof. (Proposition 1) Dropping θ, θ' from the notation, we will show that the derivative of α_σ w.r.t. σ is strictly negative.

$$\begin{aligned}
\alpha_\sigma &= \int_{-\infty}^{\infty} \min\{1, \exp(r + \sigma x - \sigma^2/2)\}\phi(x)dx \\
&= \int_{\sigma/2-r/\sigma}^{\infty} \phi(x)dx + \int_{-\infty}^{\sigma/2-r/\sigma} \exp(r + \sigma x - \sigma^2/2)\phi(x)dx \\
&= \int_{\sigma/2-r/\sigma}^{\infty} \phi(x)dx + \exp(r)\int_{-\infty}^{-\sigma/2-r/\sigma} \phi(x)dx
\end{aligned}$$

where ϕ is the probability density of the standard normal distribution. Now,

$$\begin{aligned}
\frac{d\alpha_\sigma}{d\sigma} &= -\phi(\sigma/2 - r/\sigma)(0.5 + r/\sigma^2) + \exp(r)\phi(-\sigma/2 - r/\sigma)(-0.5 + r/\sigma^2) \\
&= -\phi(\sigma/2 - r/\sigma)(0.5 + r/\sigma^2) + \phi(\sigma/2 - r/\sigma)(-0.5 + r/\sigma^2) \\
&= -\phi(\sigma/2 - r/\sigma)
\end{aligned}$$

which is negative. □